arXiv:1504.06817v1 [cs.LG] 26 Apr 2015

Analysis of Nuclear Norm Regularization for Full-rank Matrix Completion

Lijun Zhang ZHANGLJ@LAMDA.NJU.EDU.CN
National Key Laboratory for Novel Software Technology
Nanjing University, Nanjing 210023, China

Tianbao Yang TIANBAO-YANG@UIOWA.EDU
Department of Computer Science
the University of Iowa, Iowa City, IA 52242, USA

Rong Jin RONGJIN@CSE.MSU.EDU
Department of Computer Science and Engineering
Michigan State University, East Lansing, MI 48824, USA

Zhi-Hua Zhou ZHOUZH@LAMDA.NJU.EDU.CN
National Key Laboratory for Novel Software Technology
Nanjing University, Nanjing 210023, China

Abstract

In this paper, we provide a theoretical analysis of the nuclear-norm regularized least squares for full-rank matrix completion. Although similar formulations have been examined by previous studies, their results are unsatisfactory because only additive upper bounds are provided. Under the assumption that the top eigenspaces of the target matrix are incoherent, we derive a *relative* upper bound for recovering the best low-rank approximation of the unknown matrix. Our relative upper bound is tighter than previous additive bounds of other methods if the mass of the target matrix is concentrated on its top eigenspaces, and also implies perfect recovery if it is low-rank. The analysis is built upon the optimality condition of the regularized formulation and existing guarantees for low-rank matrix completion. To the best of our knowledge, this is first time such a relative bound is proved for the regularized formulation of matrix completion.

Keywords: matrix completion, nuclear norm regularization, least squares, low-rank, full-rank

1. Introduction

Matrix completion is concerned with the problem of recovering an unknown matrix from a small fraction of its entries (Candès and Tao, 2010). Recently, the problem of low-rank matrix completion has received significant interest due to theoretical advances (Candès and Recht, 2009; Keshavan et al., 2010a), as well as its applicability to a wide field of real problems, including collaborative filtering (Goldberg et al., 1992), sensor networks (Biswas et al., 2006), computer vision (Cabral et al., 2011), and machine learning (Jalali et al., 2011).

Let A be an unknown matrix of size $m \times n$, and without loss of generality, we assume $m \leq n$. The information available about A is a sampled set of entries A_{ij}, $(i, j) \in \Omega$, where Ω is a subset of the complete set of entries $[m] \times [n]$. Our goal is to recover A as precisely as possible. In a seminal work, Candès and Recht (2009) assume A is low-rank, and pose

the following nuclear norm minimization problem

$$\begin{array}{ll} \min & \|B\|_* \\ \text{s.t.} & B_{ij} = A_{ij} \ \forall (i,j) \in \Omega. \end{array} \tag{1}$$

Under the incoherence condition, Candès and Recht (2009) show that with a high probability the solution to (1) yields exact reconstruction of A, provided that a sufficiently large number of entries are observed randomly. Although the optimization problem in (1) is convex and can be formulated as a semi-definite program (Fazel et al., 2001), it is computationally expensive due to the large polynomial dependence on m and n. In the case of full-rank matrix completion, a similar nuclear norm minimization problem has been proposed. Suppose $A = Z + N$, where Z is the low-rank matrix that we want to recover, and N is the residual matrix. Candès and Plan (2010) introduce the following problem

$$\begin{array}{ll} \min & \|B\|_* \\ \text{s.t.} & \sqrt{\sum_{(i,j)\in\Omega}(B_{ij} - A_{ij})^2} \leq \delta \end{array} \tag{2}$$

where δ is an upper bound for $\sqrt{\sum_{(i,j)\in\Omega} N_{ij}^2}$. The formulation in (2) also has strong guarantees if δ is large enough, but optimization is still a challenge.

On the other hand, practitioners prefer to solve the following nuclear-norm regularized least squares problem

$$\min_{B\in\mathbb{R}^{m\times n}} \quad \frac{1}{2}\sum_{(i,j)\in\Omega}(B_{ij} - A_{ij})^2 + \lambda\|B\|_*. \tag{3}$$

for which many efficient optimization methods have been designed (Ji and Ye, 2009; Toh and Yun, 2010; Pong et al., 2010; Hsieh and Olsen, 2014). Suppose we use first-order algorithms to optimize the above problems. Due to the non-smoothness of the objective function, the convergence rates for (1) and (2) are $O(\sqrt{T})$, where T is the number of iterations (Nesterov, 2004). On the other hand, the convergence rate for (3) is $O(1/T^2)$ (Nesterov, 2013) or even linear under certain weak assumptions (Hou et al., 2013). Although (3) is computation-friendly, its recovery guarantee remains unclear. One may argue that (2) and (3) are equivalent by setting δ and λ appropriately, but the exact correspondence between them is unknown in general. We note that a similar phenomenon also happens in compressive sensing. The ℓ_1-norm minimization problem has solid theoretical guarantees (Candès and Tao, 2005; Candès, 2008), but the ℓ_1-regularized least squares is more efficient in practise (Xiao and Zhang, 2012).

To bridge the gap between practise and theory, we investigate the recovery performance of (3) theoretically. We are interested in the general case that A could be full-rank and develop theoretical guarantees for recovering the best rank-r approximation of A, denoted by A_r. In particular, we would like to measure the recovery error in terms of a relative upper bound. Let B_* be the optimal solution to (3). A relative upper bound takes the following form

$$\|B_* - A_r\|_F \leq U(r, m, n, |\Omega|)\|A - A_r\|_F$$

where $U(\cdot)$ is certain function of r, m, n and $|\Omega|$. [1] Note that such kind of bounds is very popular in compressive sensing (Cohen et al., 2009) and low-rank matrix approximation (Boutsidis et al., 2009).

Similar to previous studies, we assume the top eigenspaces of A satisfy the classical incoherence condition. Based on the celebrated results from low-rank matrix completion (Recht, 2011), we derive an upper bound for $\|B_* - A_r\|_F$, which induces a relative upper bound under favored conditions. We summarize the advantages of our results below.

- We present a general theorem that allows us to bound the recovery error of (3) for any $\lambda > 0$. In contrast, Candès and Plan (2010) only analyze the performance of (2) when $\delta \geq \sqrt{\sum_{(i,j)\in\Omega} N_{ij}^2}$.
- By choosing λ appropriately, we obtain a relative upper bound of $O(\sqrt{\frac{mnr}{|\Omega|}}\|A - A_r\|_F)$. Although similar formulations has been studied by Koltchinskii et al. (2011) and Negahban and Wainwright (2012), their bounds are additive. To the best of our knowledge, this is the *first* relative error bound for the nuclear-norm regularized formulation.
- Our relative upper bound for (3) is tighter than that for (2) developed by Candès and Plan (2010), and more general than those proved by Keshavan et al. (2010b) and Eriksson et al. (2012) under different conditions.
- Compared to the additive upper bounds of other methods (Keshavan et al., 2010b; Foygel and Srebro, 2011; Koltchinskii et al., 2011), our relative upper bound is tighter when $\|A - A_r\|_F$ is small. Notice that additive bounds never vanish even when A is low-rank.

Notations For a matrix X, we use $\|X\|_*$, $\|X\|_F$, $\|X\|$ and $\|X\|_\infty$ to denote its nuclear norm, Frobenius norm, spectral norm, and the absolute value of the largest element in magnitude, respectively,

2. Related Work

In this section, we provide a brief review of existing work.

2.1 Low-rank Matrix Completion

The mathematical study of matrix completion began with (Candès and Recht, 2009). Specifically, Candès and Recht (2009) have proved that if A obeys the incoherence condition, $|\Omega| \geq Cn^{6/5} r \log(n)$ is sufficient to ensure the convex problem in (1) succeeds with a high probability, where C is some constant that does not depends on r, m, and n. The lower bound for the size of Ω is subsequently improved to $nr\log^6(n)$ under a stronger assumption (Candès and Tao, 2010). The results presented in (Candès and Recht, 2009; Candès and Tao, 2010) are without question great breakthroughs, but their proof techniques are highly involved. In (Gross, 2011; Recht, 2011), the authors present a very elegant approach for analyzing (1), and give slightly better bounds. For example, Recht (2011) improves the bound for $|\Omega|$ to $rn\log^2(n)$ and requires the weakest assumptions on A. The simplifica-

1. By the triangle inequality, a relative upper bound for recovering A_r directly implies a relative upper bound for A.

tion of the analysis also leads to better understanding of matrix completion, and lays the foundations of the study in this paper.

In an alternative line of work, Keshavan et al. (2010a) study matrix completion using a combination of spectral techniques and manifold optimization. The proposed algorithm, which is named OPTSPACE, also achieves exact recovery if $|\Omega| \geq Cnr\max(\log(n), r)$. However, the constant C in their bound depends on many factors of A such as the aspect ratio and the condition number. After the pioneering work mentioned above, various algorithms and theories of matrix completion have been developed, including distributed matrix completion (Mackey et al., 2011), matrix completion with side information (Xu et al., 2013), 1-bit matrix completion (Cai and Zhou, 2013), coherent matrix completion (Chen et al., 2014), and universal matrix completion (Bhojanapalli and Jain, 2014), to name a few amongst many.

2.2 Full-rank Matrix Completion

We note that existing studies for full-rank matrix completion differ significantly in the assumptions they make, so their theoretical guarantees are not directly comparable. In the following, we will state existing results in the most general form, and (if possible) characterize their behaviors with respect to m, n, r, and $|\Omega|$.

Denote the optimal solution of (2) by $\widehat{B}$. Under the assumption $\delta \geq \sqrt{\sum_{(i,j)\in\Omega} N_{ij}^2}$, Theorem 7 of Candès and Plan (2010) shows

$$\|\widehat{B} - Z\|_F \leq \left(1 + m\sqrt{\frac{n}{|\Omega|}}\right)\delta$$

Let $Z = A_r$, $N = A - A_r$, and consider the optimal choice that $\delta = O\left(\sqrt{\sum_{(i,j)\in\Omega} N_{ij}^2}\right)$. The above bound becomes

$$\|\widehat{B} - A_r\|_F \leq \left(1 + m\sqrt{\frac{n}{|\Omega|}}\right)\sqrt{\sum_{(i,j)\in\Omega}(A - A_r)_{ij}^2}. \tag{4}$$

An investigation of OPTSPACE (Keshavan et al., 2010a) for full-rank matrix completion is done by Keshavan et al. (2010b). In particular, Theorem 1.1 of Keshavan et al. (2010b) implies the following additive upper bound

$$O\left(\|A_r\|_\infty m^{1/4} n^{5/4}\sqrt{\frac{r}{|\Omega|}} + \frac{mn\sqrt{r}}{|\Omega|}\|N\|\right) \tag{5}$$

where N is some matrix that depends on $A - A_r$ and Ω. We note that it is possible to derive a relative upper bound from Theorem 1.2 of Keshavan et al. (2010b), but it requires very strong assumptions about the coherence, the aspect ratio (n/m), the condition number of A_r and the r-th singular value of A. Thus, the bound derived from Theorem 1.2 of (Keshavan et al., 2010b) is much more restrictive than the bound proved here.

Foygel and Srebro (2011) study the problem of matrix completion from the view point of supervised learning. The optimization problem is formulated as least squares minimization subject to nuclear-norm or max-norm constraints. Their theoretical results follow

from generic generalization guarantees based on the Rademacher complexity. Specifically, Theorem 6 of Foygel and Srebro (2011) implies the following additive upper bound

$$O\left(\|A - A_r\|_F + n\sqrt{\frac{rm}{|\Omega|}} + \sqrt{n\|A - A_r\|_F}\left(\frac{rm}{|\Omega|}\right)^{\frac{1}{4}}\right) \tag{6}$$

where logarithmic factors are ignored.

Koltchinskii et al. (2011) have investigated a general trace regression model, which contains matrix completion as a special case. For matrix completion, they propose the following optimization problem

$$\min_{B\in\mathbb{R}^{m\times n}} \quad \frac{1}{2}\left\|B - \frac{mn}{|\Omega|}\sum_{k=1}^{|\Omega|} A_{a_k,b_k}\mathbf{e}_{a_k}\mathbf{e}_{b_k}^\top\right\|_F^2 + \lambda\|B\|_*$$

Let $\widehat{B}$ be the optimal solution to the above problem. Under certain conditions, it has been proved that with a high probability (Koltchinskii et al., 2011, Corollary 2)

$$\|\widehat{B} - A\|_F^2 + \|\widehat{B} - X\|_F^2 \leq \|X - A\|_F^2 + \frac{Cmn^2\log(n)\mathrm{rank}(X)}{|\Omega|} \tag{7}$$

for all $X \in \mathbb{R}^{m\times n}$. However, due to the presence of the second term in the upper bound, it is impossible to obtain a relative error bound.

In a recent work, Eriksson et al. (2012) consider a high-rank matrix completion problem in which the columns of A belong to a union of multiple low-rank subspaces. Under certain assumptions about the coherence as well as the geometrical arrangement of subspaces and the distribution of the columns in the subspaces, they develop a multi-step algorithm that is able to recover *each* column of A with a high probability, as long as $O(rn\log^2(m))$ entries of A are observed uniformly at random. However, the recovery guarantee of their algorithm for general full-rank matrices is unclear.

Finally, we note that a similar formulation of (3) has been studied in (Negahban and Wainwright, 2012), which differs from our paper in the following aspects.

- Negahban and Wainwright (2012) add a ℓ_∞-norm constraint to (3) and thus their optimization problem is a bit more difficult.
- Their analysis relies on the restricted strong convexity assumption, while our analysis assumes the incoherence condition.
- They derive a additive upper bound. In contrast we are able to prove a relative upper bound.

3. Our Results

We first describe the theoretical guarantees and then provide some discussions.

3.1 Theoretical Guarantees

Let $U = [\mathbf{u}_1, \ldots, \mathbf{u}_r]$ and $V = [\mathbf{v}_1, \ldots, \mathbf{v}_r]$ be two matrices that contain the first r left and right singular vectors of matrix A, respectively. Let $\mathbf{e}_i$ and $\mathbf{e}_j$ be the i-th and j-th standard basis in $\mathbb{R}^m$ and $\mathbb{R}^n$, respectively. Following the previous studies in matrix

completion (Candès and Recht, 2009; Recht, 2011), we define the coherence measure μ_0 as

$$\mu_0 = \max\left(\frac{m}{r}\max_{1\le i\le m}\|P_U\mathbf{e}_i\|^2, \frac{n}{r}\max_{1\le j\le n}\|P_V\mathbf{e}_j\|^2\right)$$

where $P_U = UU^\top$ and $P_V = VV^\top$ are two projection operators. We also define μ_1 as

$$\mu_1 = \max_{i\in[m],j\in[n]}\sqrt{\frac{mn}{r}}\left|[UV^\top]_{ij}\right|.$$

Define two projection operators $\mathcal{P}_T$ and $\mathcal{P}_{T^\perp}$ for matrices as

$$\mathcal{P}_T(Z) = P_UZ + ZP_V - P_UZP_V, \text{ and } \mathcal{P}_{T^\perp}(Z) = (I - P_U)Z(I - P_V).$$

We assume the indices are sampled uniformly with replacement, and thus Ω is a *collection* that may contain duplicate indices. The linear operator $\mathcal{R}_\Omega : \mathbb{R}^{m\times n} \mapsto \mathbb{R}^{m\times n}$ is defined as

$$\mathcal{R}_\Omega(Z) = \sum_{(i,j)\in\Omega}\langle\mathbf{e}_i\mathbf{e}_j^\top, Z\rangle\mathbf{e}_i\mathbf{e}_j^\top.$$

To simplify the notation, we define

$$\varepsilon = \|A - A_r\|_F.$$

Based on the optimality condition of B_* and the guarantees from low-rank matrix completion (Recht, 2011), we obtain the following theorem.

Theorem 1 *Assume*

$$|\Omega| \ge \max\left(114\max(\mu_0,\mu_1^2)r(m+n)\beta\log^2(2n), \frac{8mn\|A - A_r\|_\infty^2}{3\|A - A_r\|_F^2}\beta\log(n)\right) \tag{8}$$

for some $\beta > 1$, and $n \ge 5$. With a probability at least $1 - 6\log(n)(m+n)^{2-2\beta} - n^{2-2\beta^{1/2}} - n^{-\beta}$, we have

$$\|\mathcal{P}_{T^\perp}(B_*)\|_F \le \|\mathcal{P}_{T^\perp}(B_*)\|_* \le \frac{8|\Omega|\varepsilon^2}{mn\lambda} + \frac{3mnr\log(2n)\lambda}{|\Omega|},$$

$$\|\mathcal{P}_T(A_r - B_*)\|_F \le 4\varepsilon + \frac{2mn\lambda}{|\Omega|}\sqrt{3r\log(2n)} + 64\log(n)\sqrt{\frac{mn\beta}{6|\Omega|}}\|\mathcal{P}_{T^\perp}(B_*)\|_F.$$

As can be seen, our upper bound is valid for any $\lambda > 0$. In contrast, the upper bound for (2) in (Candès and Plan, 2010) is limited to the case $\delta \ge \sqrt{\langle\mathcal{R}_\Omega(A - A_r, A - A_r)\rangle}$.

By choosing λ to minimize the upper bounds in the above theorem, we obtain the following corollary.

Corollary 2 *Under the condition in Theorem 1. Set*

$$\lambda = \frac{2|\Omega|\varepsilon}{mn}\sqrt{\frac{2}{3r\log(2n)}}.$$

With a probability at least $1 - 6\log(n)(m+n)^{2-2\beta} - n^{2-2\beta^{1/2}} - n^{-\beta}$, *we have*

$$\|\mathcal{P}_{T^\perp}(B_*)\|_F \leq 4\sqrt{6r\log(2n)}\varepsilon,$$
$$\|\mathcal{P}_T(A_r - B_*)\|_F \leq \left(10 + 256\sqrt{\frac{mnr\log^3(2n)\beta}{|\Omega|}}\right)\varepsilon,$$

and thus

$$\|A_r - B_*\|_F \leq O\left(\sqrt{r\log(n)} + \sqrt{\frac{mnr\log^3(n)}{|\Omega|}}\right)\varepsilon.$$

3.2 Discussions

Comparisons Let's consider the most practical scenario $|\Omega| \leq mn$, and for simplicity ignore logarithmic factors. In this case, our relative bound becomes

$$\|A_r - B_*\|_F \leq O\left(\sqrt{\frac{mnr}{|\Omega|}}\right)\varepsilon.$$

The most comparable result is the relative upper bound derived by Candès and Plan (2010), because they also rely on the incoherence condition. Using Lemma 2 derived later in this paper to simplify (4) in Section 2.2, we obtain the following relative upper bound

$$\|A_r - \widehat{B}\|_F \leq O\left(\sqrt{m}\|A - A_r\|_F\right) = O\left(\sqrt{m}\varepsilon\right)$$

which is always worse than our bound since $|\Omega| \geq Cnr$ for some constant C. Compared to the relative upper bounds of Keshavan et al. (2010b) and Eriksson et al. (2012), our result is more general. Because their bounds only hold for a very restricted class of matrices.

Our relative bound is tighter than the additive bound in (5) derived by Keshavan et al. (2010b), if

$$\varepsilon \leq O\left(\frac{n^{3/4}}{m^{1/4}}\right),$$

and also tighter than the additive bound in (6) derived by Foygel and Srebro (2011), if $\varepsilon \leq O(\sqrt{n})$. Let $X = A_r$, the additive bound in (7) derived by Koltchinskii et al. (2011) implies

$$\left\|\widehat{B} - A_r\right\|_F \leq \varepsilon + O\left(\sqrt{\frac{mn^2r}{|\Omega|}}\right)$$

which is worse than our bound if $\varepsilon \leq O(\sqrt{n})$.

Low-rank Case In the special case that A is a rank-r matrix, we should interpret Corollary 2 as describing the limiting behavior of (3). It implies the recovery error will approach 0 as $\lambda \to 0$. That is, we can set λ to be an arbitrarily small constant and the recovery error is also arbitrarily small.

Assumptions The only assumption that we make is about the cardinality of Ω in (8). The first lower bound is essential since it is the necessary condition for us to utilize the theoretical guarantees developed for low-rank matrix completion (Recht, 2011). In contrast,

the second lower bound is not obligatory. It is just used to facilitate a simple application of Bernstein's inequality in Lemma 2. Without the second lower bound, we still get a relative upper bound, but at a larger order. For details, please refer to Section 4.4.

Extensions Although the current result is built upon the result in (Recht, 2011), which in turn requires the incoherence assumption, it can be easily extended to support other assumptions for matrix completion. What we need is to replace Theorem 3 in Section 4.1 with the corresponding theorem derived under other assumptions. We will still get a relative upper bound, possibly at different order. For example, if we use the theorems in (Bhojanapalli and Jain, 2014), our bound becomes a *universal* guarantee for full-matrix completion. We leave the extension of our analysis to other assumptions as a future work.

4. Analysis

We present the proof of Theorem 1 in this section.

4.1 Sketch of the Proof

As we mentioned before, our analysis is built upon the existing theoretical guarantees for low-rank matrix completion, which is summarized below (Recht, 2011).

Theorem 3 *Suppose*

$$|\Omega| \geq 32 \max(\mu_0, \mu_1^2) r(m+n)\beta \log^2(2n) \tag{9}$$

for some $\beta > 1$. Then, with a probability at least $1 - 6\log(n)(m+n)^{2-2\beta} - n^{2-2\beta^{1/2}}$, the following statements are true:

- $$\left\| \frac{mn}{|\Omega|} \mathcal{P}_T \mathcal{R}_\Omega \mathcal{P}_T - \mathcal{P}_T \right\| \leq \frac{1}{2}. \tag{10}$$
- $$\|\mathcal{R}_\Omega\| \leq \frac{8}{3}\sqrt{\beta}\log(n). \tag{11}$$
- *There exists a $Y \in \mathbb{R}^{m\times n}$ in the range of $\mathcal{R}_\Omega$ such that*

$$\left\| \mathcal{P}_T(Y) - UV^\top \right\|_F \leq \sqrt{\frac{r}{2n}}, \tag{12}$$

$$\left\| \mathcal{P}_{T^\perp(Y)} \right\| \leq \frac{1}{2}, \tag{13}$$

$$|\langle Y, A\rangle| \leq \sqrt{\frac{3mnr\log(2n)}{8|\Omega|}} \sqrt{\langle \mathcal{R}_\Omega(A), A\rangle}, \tag{14}$$

for all $A \in \mathbb{R}^{m\times n}$.

The first part of above theorem contains concentration inequalities for the random linear operator $\mathcal{P}_T\mathcal{R}_\Omega\mathcal{P}_T$ and $\mathcal{R}_\Omega$, and the second part describes some important properties of a special matrix Y, which is used as an (approximate) dual certificate of (1).

Next, we will examine the optimality of B_* based on techniques from convex analysis, leading to the following theorem.

Theorem 4 *Let B_* be the optimal solution to (3), we have*

$$\lambda\langle B_* - A_r, UV^\top\rangle + \lambda\|\mathcal{P}_{T^\perp}(B_*)\|_* \leq \langle \mathcal{R}_\Omega(B_* - A), A_r - B_*\rangle. \tag{15}$$

Based on Theorems 3 and 4, we are ready to prove the main results. However, the analysis is a bit lengthy, so we split it into two parts, and will first show the following intermediate theorem.

Theorem 5 *Under the condition in Theorem 1. With a probability at least* $1-6\log(n)(m+n)^{2-2\beta} - n^{2-2\beta^{1/2}} - n^{-\beta}$*, we have*

$$\frac{1}{2}\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle + \frac{\lambda}{2}\left\|\mathcal{P}_{T^\perp}(B_*)\right\|_* \leq \lambda\sqrt{\frac{r}{2n}}\|\mathcal{P}_T(A_r - B_*)\|_F + \Gamma \tag{16}$$

where

$$\Gamma = \frac{2|\Omega|\varepsilon^2}{mn} + \frac{3mnr\log(2n)\lambda^2}{8|\Omega|}. \tag{17}$$

Before going to the detail, we introduce a lemma that will be used throughout the analysis. Since Ω may contain duplicate indices, $\langle \mathcal{R}_\Omega(A), A\rangle \neq \|\mathcal{R}_\Omega(A)\|_F^2$ in general. We use the following lemma to take care of this issue.

Lemma 1

$$\langle \mathcal{R}_\Omega(A), A\rangle \leq \|\mathcal{R}_\Omega(A)\|_F^2, \tag{18}$$

$$|\langle \mathcal{R}_\Omega(A), B\rangle| \leq \sqrt{\langle \mathcal{R}_\Omega(A), A\rangle}\sqrt{\langle \mathcal{R}_\Omega(B), B\rangle} \tag{19}$$

for all $A, B \in \mathbb{R}^{m\times n}$.

4.2 Proof of Theorem 3

Except for the last inequality in (14), all the others can be found directly from the proof of Theorem 2 in (Recht, 2011). Thus, we only provide the derivation of (14). Following the construction in (Recht, 2011), we partition Ω into p partitions of size q. By assumption, we can choose

$$q \geq \frac{128}{3}\max(\mu_0, \mu_1^2) r(m+n)\beta\log(m+n)$$

such that

$$p = \frac{\Omega}{q} = \frac{3}{4}\log 2n.$$

Let Ω_j denote the set of indices corresponding to the j-th partition. We define $W_0 = UV^\top$,

$$\begin{aligned} Y_k =& \frac{mn}{q}\sum_{j=1}^{k}\mathcal{R}_{\Omega_j}(W_{j-1}), \\ W_k =& UV^\top - \mathcal{P}_T(Y_k), \end{aligned}$$

for $k = 1, \ldots, q$. Then, we set $Y = Y_p$. It has been proved that with a probability at least $1 - 6\log(n)(m+n)^{2-2\beta} - n^{2-2\beta^{1/2}}$,

$$\|W_k\|_F \leq 2^{-k}\sqrt{r}, \tag{20}$$

$$\left\| \frac{mn}{q}\mathcal{P}_T\mathcal{R}_{\Omega_k}\mathcal{P}_T - \mathcal{P}_T \right\| \leq \frac{1}{2}, \tag{21}$$

for $k = 1, \ldots, q$.

Since $W_j = \mathcal{P}_T(W_j)$, we have

$$\begin{aligned} \frac{mn}{q}\langle \mathcal{R}_{\Omega_j}(W_j), W_j\rangle &= \left\langle W_j, \frac{mn}{q}\mathcal{P}_T\mathcal{R}_{\Omega_j}\mathcal{P}_T(W_j)\right\rangle \\ \overset{(21)}{\leq} \frac{3}{2}\left\|\mathcal{P}_T(W_j)\right\|_F^2 &= \frac{3}{2}\|W_j\|_F^2 \overset{(20)}{\leq} \frac{3r}{2}4^{-j}. \end{aligned} \tag{22}$$

Then,

$$\begin{aligned} |\langle Y, A\rangle| \leq & \frac{mn}{q}\sum_{j=1}^{p}\left|\langle \mathcal{R}_{\Omega_j}(W_{j-1}), A\rangle\right| \\ \overset{(19)}{\leq} & \frac{mn}{q}\sum_{j=1}^{p}\sqrt{\langle \mathcal{R}_{\Omega_j}(W_{j-1}), W_{j-1}\rangle}\sqrt{\langle \mathcal{R}_{\Omega_j}(A), A\rangle} \\ \leq & \frac{mn}{q}\sqrt{\sum_{j=1}^{p}\langle \mathcal{R}_{\Omega_j}(W_{j-1}), W_{j-1}\rangle}\sqrt{\sum_{j=1}^{p}\langle \mathcal{R}_{\Omega_j}(A), A\rangle} \\ = & \sqrt{\frac{mn}{q}}\sqrt{\sum_{j=1}^{p}\left\langle \frac{mn}{q}\mathcal{R}_{\Omega_j}(W_{j-1}), W_{j-1}\right\rangle}\sqrt{\langle \mathcal{R}_{\Omega}(A), A\rangle} \\ \overset{(22)}{\leq} & \sqrt{\frac{mn}{q}}\sqrt{\langle \mathcal{R}_{\Omega}(A), A\rangle}\sqrt{\frac{3r}{2}\sum_{j=1}^{p}4^{-j}} \\ \leq & \sqrt{\frac{mnr}{2q}}\sqrt{\langle \mathcal{R}_{\Omega}(A), A\rangle} = \sqrt{\frac{mnpr}{2|\Omega|}}\sqrt{\langle \mathcal{R}_{\Omega}(A), A\rangle}. \end{aligned}$$

4.3 Proof of Theorem 4

Since B_* is the optimal solution to (3), we have

$$\langle \mathcal{R}_{\Omega}(B_* - A) + \lambda E, A_r - B_*\rangle \geq 0 \tag{23}$$

where $E \in \partial\|B_*\|_*$ is *certain* subgradient of $\|\cdot\|_*$ evaluated at B_*. Let $F \in \partial\|A_r\|_*$ be *any* subgradient of $\|\cdot\|_*$ evaluated at A_r. From the property of convexity, we have

$$\langle B_* - A_r, E - F\rangle \geq 0. \tag{24}$$

From (23) and (24), we get

$$\langle \mathcal{R}_{\Omega}(B_* - A) + \lambda F, A_r - B_*\rangle \geq 0. \tag{25}$$

Next, we consider bounding $\lambda\langle F, A_r - B_*\rangle$. From previous studies (Candès and Recht, 2009), we know that the set of subgradients of $\|A_r\|_*$ takes the following form:

$$\partial\|A_r\|_* = \left\{UV^\top + W : W \in \mathbb{R}^{m\times n}, U^\top W = 0, WV = 0, \|W\| \le 1\right\}.$$

Thus, we can choose

$$F = UV^\top + \mathcal{P}_{T^\perp}(N),$$

where $N = \text{argmax}_{\|X\|\le 1}\langle \mathcal{P}_{T^\perp}(B_*), X\rangle$. Then, it is easy to verify that

$$\begin{aligned}\langle B_* - A_r, F\rangle =& \langle B_* - A_r, UV^\top\rangle + \langle B_* - A_r, \mathcal{P}_{T^\perp}(N)\rangle \\ =& \langle B_* - A_r, UV^\top\rangle + \|\mathcal{P}_{T^\perp}(B_*)\|_*.\end{aligned} \tag{26}$$

We complete the proof by combining (25) and (26).

4.4 Proof of Theorem 5

We first introduce a lemma that will be used later.

Lemma 2 *Suppose*

$$|\Omega| \ge \frac{8mn\|A - A_r\|_\infty^2}{3\|A - A_r\|_F^2}\beta\log(n) \tag{27}$$

for some $\beta > 1$. Then, with a probability at least $1 - n^{-\beta}$, we have

$$\sqrt{\langle \mathcal{R}_\Omega(A - A_r), A - A_r\rangle} \le \varepsilon\sqrt{\frac{2|\Omega|}{mn}}. \tag{28}$$

If the condition in (27) does not hold, we can use the following inequality

$$\sqrt{\langle \mathcal{R}_\Omega(A - A_r), A - A_r\rangle} \overset{(18)}{\le} \|\mathcal{R}_\Omega(A - A_r)\|_F \overset{(11)}{\le} \frac{8}{3}\sqrt{\beta}\log(n)\varepsilon$$

and will obtain a similar bound but at a larger order.

We continue the proof by lower bounding $\langle B_* - A_r, UV^\top\rangle$ in (15) of Theorem 4. To this end, we need the matrix Y given in Theorem 3.

$$\begin{aligned}\langle B_* - A_r, UV^\top\rangle =& \langle B_* - A_r, UV^\top - Y\rangle + \langle B_* - A_r, Y\rangle \\ =& \langle B_* - A_r, UV^\top - \mathcal{P}_T(Y)\rangle + \langle A_r - B_*, \mathcal{P}_{T^\perp}(Y)\rangle + \langle B_* - A_r, Y\rangle.\end{aligned}$$

Next, we bound the last three terms by utilizing the conclusions in Theorem 3.

$$\begin{aligned}&\langle B_* - A_r, UV^\top - \mathcal{P}_T(Y)\rangle = \langle \mathcal{P}_T(B_* - A_r), UV^\top - \mathcal{P}_T(Y)\rangle \\ \ge& -\|\mathcal{P}_T(B_* - A_r)\|_F\|UV^\top - \mathcal{P}_T(Y)\|_F \overset{(12)}{\ge} -\sqrt{\frac{r}{2n}}\|\mathcal{P}_T(A_r - B_*)\|_F.\end{aligned}$$

$$\begin{aligned}&\langle A_r - B_*, \mathcal{P}_{T^\perp}(Y)\rangle = \langle \mathcal{P}_{T^\perp}(A_r - B_*), \mathcal{P}_{T^\perp}(Y)\rangle \\ =& \langle \mathcal{P}_{T^\perp}(-B_*), \mathcal{P}_{T^\perp}(Y)\rangle \ge -\|\mathcal{P}_{T^\perp}(B_*)\|_*\|\mathcal{P}_{T^\perp}(Y)\| \overset{(13)}{\ge} -\frac{1}{2}\|\mathcal{P}_{T^\perp}(B_*)\|_*.\end{aligned}$$

$$\langle B_* - A_r, Y\rangle \overset{(14)}{\geq} -\sqrt{\frac{3mnr\log(2n)}{8|\Omega|}}\sqrt{\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle}.$$

Putting the above inequalities together, we have

$$\begin{aligned}\langle B_* - A_r, UV^\top\rangle \geq &-\sqrt{\frac{r}{2n}}\|\mathcal{P}_T(A_r - B_*)\|_F - \frac{1}{2}\left\|\mathcal{P}_{T^\perp}(B_*)\right\|_* \\ &-\sqrt{\frac{3mnr\log(2n)}{8|\Omega|}}\sqrt{\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle}.\end{aligned} \tag{29}$$

Substituting (29) into (15) and rearranging, we get

$$\begin{aligned}&\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle + \frac{\lambda}{2}\left\|\mathcal{P}_{T^\perp}(B_*)\right\|_* \\ \leq&\langle \mathcal{R}_\Omega(A_r - A), A_r - B_*\rangle + \lambda\sqrt{\frac{r}{2n}}\|\mathcal{P}_T(A_r - B_*)\|_F \\ &+ \lambda\sqrt{\frac{3mnr\log(2n)}{8|\Omega|}}\sqrt{\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle}.\end{aligned} \tag{30}$$

Furthermore, we have

$$\begin{aligned}&\langle \mathcal{R}_\Omega(A_r - A), A_r - B_*\rangle \\ \overset{(19)}{\leq}&\sqrt{\langle \mathcal{R}_\Omega(A - A_r), A - A_r\rangle}\sqrt{\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle} \\ \overset{(28)}{\leq}&\varepsilon\sqrt{\frac{2|\Omega|}{mn}}\sqrt{\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle}\end{aligned} \tag{31}$$

where Lemma 2 is used in the last inequality.

From (30) and (31), we have

$$\begin{aligned}&\mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle + \frac{\lambda}{2}\left\|\mathcal{P}_{T^\perp}(B_*)\right\|_* \\ \leq&\lambda\sqrt{\frac{r}{2n}}\|\mathcal{P}_T(A_r - B_*)\|_F + \varepsilon\sqrt{\frac{2|\Omega|}{mn}}\sqrt{\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle} \\ &+ \lambda\sqrt{\frac{3mnr\log(2n)}{8|\Omega|}}\sqrt{\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle}.\end{aligned}$$

We complete the proof by using the basic inequality

$$\frac{1}{4}\alpha^2 - \alpha\beta + \beta^2 \geq 0$$

with $\alpha = \sqrt{\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle}$ and $\beta \in \{\varepsilon\sqrt{\frac{2|\Omega|}{mn}}, \lambda\sqrt{\frac{3mnr\log(2n)}{8|\Omega|}}\}$.

4.5 Proof of Theorem 1

We first explain the two lower bounds of $|\Omega|$ in (8). The first one is due to Theorem 3, but we use a larger constant (114 instead of 32) to ensure

$$\frac{8\log(n)}{3}\sqrt{\frac{rm\beta}{|\Omega|}} \leq \frac{1}{4} \tag{32}$$

which is used later. The second one is due to Lemma 2.

4.5.1 Upper Bound for $\|\mathcal{P}_{T^\perp}(B_*)\|_F$

We upper bound $\|\mathcal{P}_T(A_r - B_*)\|_F^2$ in (16) of Theorem 5 by

$$\|\mathcal{P}_T(A_r - B_*)\|_F^2 = \langle \mathcal{P}_T(A_r - B_*), A_r - B_*\rangle \overset{(10)}{\leq} 2\frac{mn}{|\Omega|}\langle \mathcal{P}_T\mathcal{R}_\Omega\mathcal{P}_T(A_r - B_*), A_r - B_*\rangle.$$

Plugging the above inequality in (16), we have

$$\frac{1}{2}\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle + \frac{\lambda}{2}\left\|\mathcal{P}_{T^\perp}(B_*)\right\|_* \leq \lambda\sqrt{\frac{rm}{|\Omega|}}\sqrt{\langle \mathcal{P}_T\mathcal{R}_\Omega\mathcal{P}_T(A_r - B_*), A_r - B_*\rangle} + \Gamma. \tag{33}$$

Since $\mathcal{P}_T + \mathcal{P}_{T^\perp} = I$, we have

$$\begin{aligned}
&\frac{1}{2}\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle \\
=&\frac{1}{2}\underbrace{\langle \mathcal{P}_T\mathcal{R}_\Omega\mathcal{P}_T(A_r - B_*), A_r - B_*\rangle}_{:=\Theta^2} + \frac{1}{2}\underbrace{\langle \mathcal{P}_{T^\perp}\mathcal{R}_\Omega\mathcal{P}_{T^\perp}(B_*), B_*\rangle}_{:=\Lambda^2} \\
&- \langle \mathcal{R}_\Omega(\mathcal{P}_T(A_r - B_*)), \mathcal{P}_{T^\perp}(B_*)\rangle \\
\overset{(19)}{\geq}& \frac{1}{2}\Theta^2 + \frac{1}{2}\Lambda^2 - \Theta\Lambda = \frac{1}{2}(\Theta - \Lambda)^2.
\end{aligned} \tag{34}$$

Substituting (34) into (33), we have

$$\frac{1}{2}(\Theta - \Lambda)^2 + \frac{\lambda}{2}\left\|\mathcal{P}_{T^\perp}(B_*)\right\|_* \leq \lambda\sqrt{\frac{rm}{|\Omega|}}\Theta + \Gamma.$$

Combining with the fact

$$\frac{1}{2}(\Theta - \Lambda)^2 - \lambda\sqrt{\frac{rm}{|\Omega|}}\Theta + \lambda\sqrt{\frac{rm}{|\Omega|}}\Lambda + \frac{rm\lambda^2}{2|\Omega|} \geq 0$$

we have

$$\begin{aligned}\frac{\lambda}{2}\left\|\mathcal{P}_{T^\perp}(B_*)\right\|_* \leq & \lambda\sqrt{\frac{rm}{|\Omega|}}\sqrt{\langle\mathcal{P}_{T^\perp}\mathcal{R}_\Omega\mathcal{P}_{T^\perp}(B_*),B_*\rangle}+\frac{rm\lambda^2}{2|\Omega|}+\Gamma\\ \overset{(18)}{\leq} & \lambda\sqrt{\frac{rm}{|\Omega|}}\|\mathcal{R}_\Omega\mathcal{P}_{T^\perp}(B_*)\|_F+\frac{rm\lambda^2}{2|\Omega|}+\Gamma\\ \overset{(11)}{\leq} & \frac{8\lambda\log(n)}{3}\sqrt{\frac{rm\beta}{|\Omega|}}\|\mathcal{P}_{T^\perp}(B_*)\|_F+\frac{rm\lambda^2}{2|\Omega|}+\Gamma\\ \overset{(32)}{\leq} & \frac{\lambda}{4}\|\mathcal{P}_{T^\perp}(B_*)\|_F+\frac{rm\lambda^2}{2|\Omega|}+\Gamma\\ \overset{(17)}{\leq} & \frac{\lambda}{4}\|\mathcal{P}_{T^\perp}(B_*)\|_*+\frac{2|\Omega|\varepsilon^2}{mn}+\frac{3mnr\log(2n)\lambda^2}{4|\Omega|}\end{aligned}$$

which implies

$$\left\|\mathcal{P}_{T^\perp}(B_*)\right\|_* \leq \frac{8|\Omega|\varepsilon^2}{mn\lambda}+\frac{3mnr\log(2n)\lambda}{|\Omega|}. \tag{35}$$

4.5.2 Upper Bound for $\|\mathcal{P}_T(A_r-B_*)\|_F$

Similar to (34), we have

$$\begin{aligned}&\frac{1}{2}\langle\mathcal{R}_\Omega(A_r-B_*),A_r-B_*\rangle\\ \overset{(19)}{\geq} & \frac{1}{2}\langle\mathcal{P}_T\mathcal{R}_\Omega\mathcal{P}_T(A_r-B_*),A_r-B_*\rangle+\frac{1}{2}\Lambda^2-\sqrt{\langle\mathcal{P}_T\mathcal{R}_\Omega\mathcal{P}_T(A_r-B_*),A_r-B_*\rangle}\Lambda\\ \overset{(10)}{\geq} & \frac{|\Omega|}{4mn}\|\mathcal{P}_T(A_r-B_*)\|_F^2+\frac{1}{2}\Lambda^2-\sqrt{\frac{3|\Omega|}{2mn}}\|\mathcal{P}_T(A_r-B_*)\|_F\Lambda\end{aligned}$$

where

$$\Lambda=\sqrt{\langle\mathcal{P}_{T^\perp}\mathcal{R}_\Omega\mathcal{P}_{T^\perp}(B_*),B_*\rangle}\overset{(18)}{\leq}\|\mathcal{R}_\Omega\mathcal{P}_{T^\perp}(B_*)\|_F\overset{(11)}{\leq}\frac{8}{3}\sqrt{\beta}\log(n)\|\mathcal{P}_{T^\perp}(B_*)\|_F.$$

By plugging the above inequalities into (16), we have

$$\begin{aligned}&\frac{|\Omega|}{4mn}\|\mathcal{P}_T(A_r-B_*)\|_F^2+\frac{1}{2}\Lambda^2+\frac{\lambda}{2}\left\|\mathcal{P}_{T^\perp}(B_*)\right\|_*\\ \leq & \lambda\sqrt{\frac{r}{2n}}\|\mathcal{P}_T(A_r-B_*)\|_F+\Gamma+8\log(n)\sqrt{\frac{|\Omega|\beta}{6mn}}\|\mathcal{P}_{T^\perp}(B_*)\|_F\|\mathcal{P}_T(A_r-B_*)\|_F\end{aligned}$$

and thus

$$\begin{aligned}\|\mathcal{P}_T(A_r-B_*)\|_F^2 \overset{(17)}{\leq} & \frac{2m\lambda\sqrt{2rn}}{|\Omega|}\|\mathcal{P}_T(A_r-B_*)\|_F+8\varepsilon^2+\frac{3rm^2n^2\lambda^2\log(2n)}{2|\Omega|^2}\\ & +32\log(n)\sqrt{\frac{mn\beta}{6|\Omega|}}\|\mathcal{P}_{T^\perp}(B_*)\|_F\|\mathcal{P}_T(A_r-B_*)\|_F.\end{aligned}$$

Recall that

$$x^2 \leq bx + c \Rightarrow x \leq 2b + \sqrt{2c}.$$

Thus, we have

$$\|\mathcal{P}_T(A_r - B_*)\|_F \leq \frac{4m\lambda\sqrt{2rn}}{|\Omega|} + 4\varepsilon + \frac{mn\lambda}{|\Omega|}\sqrt{3r\log(2n)} + 64\log(n)\sqrt{\frac{mn\beta}{6|\Omega|}}\|\mathcal{P}_{T^\perp}(B_*)\|_F.$$

We complete the proof by noticing

$$4\sqrt{2n} \leq n\sqrt{3\log(2n)},\ \forall n \geq 5.$$

4.6 Proof of Lemma 1

Denote the number of unique indices in Ω be u, and let $\Theta = \{(a_k, b_k)\}_{k=1}^u$ be a *set* that contains all the unique indices in Ω. Let t_k denote the times that (a_k, b_k) appears in Ω. Then, we have

$$\langle \mathcal{R}_\Omega(A), A\rangle = \sum_{k=1}^u t_k A_{a_k b_k}^2 \leq \sum_{k=1}^u t_k^2 A_{a_k b_k}^2 = \|\mathcal{R}_\Omega(A)\|_F^2.$$

To show (19), we have

$$|\langle \mathcal{R}_\Omega(A), B\rangle| = \left|\sum_{(i,j)\in\Omega} A_{ij}B_{ij}\right| \leq \sqrt{\sum_{(i,j)\in\Omega} A_{ij}^2}\sqrt{\sum_{(i,j)\in\Omega} B_{ij}^2} = \sqrt{\langle \mathcal{R}_\Omega(A), A\rangle}\sqrt{\langle \mathcal{R}_\Omega(B), B\rangle}$$

where the second line follows from Cauchy–Schwarz inequality.

4.7 Proof of Lemma 2

For each index $(a_k, b_k) \in \Omega$, we define a random variable

$$\xi_k = \langle \mathbf{e}_{a_k}\mathbf{e}_{b_k}^\top, A - A_r\rangle^2 - \frac{1}{mn}\|A - A_r\|_F^2.$$

Then, it is easy to verify that

$$\begin{aligned}
\mathrm{E}[\xi_k] =& 0,\\
|\xi_k| =& \left|\langle \mathbf{e}_{a_k}\mathbf{e}_{b_k}^\top, A - A_r\rangle^2 - \frac{1}{mn}\|A - A_r\|_F^2\right|\\
\leq& \max\left(\langle \mathbf{e}_{a_k}\mathbf{e}_{b_k}^\top, A - A_r\rangle^2, \frac{1}{mn}\|A - A_r\|_F^2\right) \leq \|A - A_r\|_\infty^2,\\
\mathrm{E}[\xi_k^2] =& \mathrm{E}\left[\langle \mathbf{e}_{a_k}\mathbf{e}_{b_k}^\top, A - A_r\rangle^4\right] - \frac{1}{m^2n^2}\|A - A_r\|_F^4\\
\leq& \frac{1}{mn}\sum_{i,j}[A - A_r]_{ij}^4 \leq \frac{1}{mn}\|A - A_r\|_\infty^2\|A - A_r\|_F^2.
\end{aligned}$$

From Bernstein's inequality, we have

$$\begin{aligned}
&\mathrm{P}\left[\langle \mathcal{R}_\Omega(A-A_r), A-A_r\rangle \geq 2\frac{|\Omega|}{mn}\|A-A_r\|_F^2\right] \\
=&\mathrm{P}\left[\sum_{k=1}^{|\Omega|}\xi_k \geq \frac{|\Omega|}{mn}\|A-A_r\|_F^2\right] \leq \exp\left(-\frac{3|\Omega|}{8\|A-A_r\|_\infty^2}\frac{1}{mn}\|A-A_r\|_F^2\right) \overset{(27)}{\leq} n^{-\beta}.
\end{aligned}$$

5. Conclusion and Future Work

In this paper, we provide a theoretical analysis of the nuclear-norm regularized formulation in (3) for matrix completion. Assuming the top eigenspaces are incoherent, a relative upper bound is derived. An extensive comparisons demonstrate that our bound is tighter than previous results under favored conditions.

In certain real-world scenarios, we may further assume the observations are corrupted by noise. While Theorem 1 only addresses the noise-free case, we can immediately extend our results to the noisy case. Let N be the matrix of noise. We just need to add $\langle \mathcal{R}_\Omega(N), A_r - B_*\rangle$ to the right hand side of (15), and upper bound it by $\sqrt{\langle \mathcal{R}_\Omega(N), N\rangle}\sqrt{\langle \mathcal{R}_\Omega(A_r - B_*), A_r - B_*\rangle}$. Then, we redefine Γ to include $\sqrt{\langle \mathcal{R}_\Omega(N), N\rangle}$, and the rest of the proof is almost the same.

One limitation of the current analysis is that the optimal value of λ depends on $\|A - A_r\|_F$, which is usually unknown. The same problem is also suffered by the nuclear norm minimization problem in (2). We will investigate how to estimate λ in the future.

References

Srinadh Bhojanapalli and Prateek Jain. Universal matrix completion. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.

Pratik Biswas, Tzu-Chen Lian, Ta-Chung Wang, and Yinyu Ye. Semidefinite programming based algorithms for sensor network localization. *ACM Transactions on Sensor Networks*, 2(2):188–220, 2006.

Christos Boutsidis, Michael W. Mahoney, and Petros Drineas. An improved approximation algorithm for the column subset selection problem. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 968–977, 2009.

Ricardo S. Cabral, Fernando Torre, Joao P. Costeira, and Alexandre Bernardino. Matrix completion for multi-label image classification. In *Advances in Neural Information Processing Systems 24*, pages 190–198, 2011.

Tony Cai and Wen-Xin Zhou. A max-norm constrained minimization approach to 1-bit matrix completion. *Journal of Machine Learning Research*, 14:3619–3647, 2013.

Emmanuel J. Candès. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathematique*, 346(9–10):589–592, 2008.

Emmanuel J. Candès and Yaniv Plan. Matrix completion with noise. *Proceedings of the IEEE*, 98(6):925–936, 2010.

Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

Emmanuel J. Candès and Terence Tao. Decoding by linear programming. *Information Theory, IEEE Transactions on*, 51(12):4203–4215, 2005.

Emmanuel J. Candès and Terence Tao. The power of convex relaxation: Near-optimal matrix completion. *IEEE Transactions on Information Theory*, 56(5):2053–2080, 2010.

Yudong Chen, Srinadh Bhojanapalli, Sujay Sanghavi, and Rachel Ward. Coherent matrix completion. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.

Albert Cohen, Wolfgang Dahmen, and Ronald DeVore. Compressed sensing and best k-term approximation. *Journal of the American Mathematical Society*, 22(1):211C–231, 2009.

Brian Eriksson, Laura Balzano, and Robert Nowak. High-rank matrix completion. In *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics*, pages 373–381, 2012.

Maryam Fazel, Haitham Hindi, and Stephen P. Boyd. A rank minimization heuristic with application to minimum order system approximation. In *Proceedings American Control Conference*, pages 4734–4739, 2001.

Rina Foygel and Nathan Srebro. Concentration-based guarantees for low-rank matrix reconstruction. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 315–340, 2011.

David Goldberg, David Nichols, Brian M. Oki, and Douglas Terry. Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12):61–70, 1992.

David Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Transactions on Information Theory,*, 57(3):1548–1566, 2011.

Ke Hou, Zirui Zhou, Anthony Man-Cho So, and Zhi-Quan Luo. On the linear convergence of the proximal gradient method for trace norm regularization. In *Advances in Neural Information Processing Systems 26*, pages 710–718, 2013.

Cho-Jui Hsieh and Peder Olsen. Nuclear norm minimization via active subspace selection. In *Proceedings of the 31st International Conference on Machine Learning*, 2014.

Ali Jalali, Yudong Chen, Sujay Sanghavi, and Huan Xu. Clustering partially observed graphs via convex optimization. In *Proceedings of the 28th International Conference on Machine Learning*, pages 1001–1008, 2011.

Shuiwang Ji and Jieping Ye. An accelerated gradient method for trace norm minimization. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 457–464, 2009.

Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56(6):2980–2998, 2010a.

Raghunandan H. Keshavan, Andrea Montanari, and Sewoong Oh. Matrix completion from noisy entries. *Journal of Machine Learning Research*, 11:2057–2078, 2010b.

Vladimir Koltchinskii, Karim Lounici, and Alexandre B. Tsybakov. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.

Lester W. Mackey, Michael I. Jordan, and Ameet Talwalkar. Divide-and-conquer matrix factorization. In *Advances in Neural Information Processing Systems 24*, pages 1134–1142, 2011.

Sahand Negahban and Martin J. Wainwright. Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, 13:1665–1697, 2012.

Yu. Nesterov. Gradient methods for minimizing composite functions. *Mathematical Programming*, 140(1):125–161, 2013.

Yurii Nesterov. *Introductory lectures on convex optimization: a basic course*, volume 87 of *Applied optimization*. Kluwer Academic Publishers, 2004.

Ting Kei Pong, Paul Tseng, Shuiwang Ji, and Jieping Ye. Trace norm regularization: Reformulations, algorithms, and multi-task learning. *SIAM Journal on Optimization*, 20 (6):3465–3489, 2010.

Benjamin Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12:3413–3430, 2011.

Kim-Chuan Toh and Sangwoon Yun. An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. *Pacific Journal of Optimization*, 6(3):615–640, 2010.

Lin Xiao and Tong Zhang. A proximal-gradient homotopy method for the ℓ_1-regularized least-squares problem. In *Proceedings of the 29th International Conference on Machine Learning*, pages 839–846, 2012.

Miao Xu, Rong Jin, and Zhi-Hua Zhou. Speedup matrix completion with side information: Application to multi-label learning. In *Advances in Neural Information Processing Systems 26*, pages 2301–2309, 2013.